
Database Management Systems Technology Trends

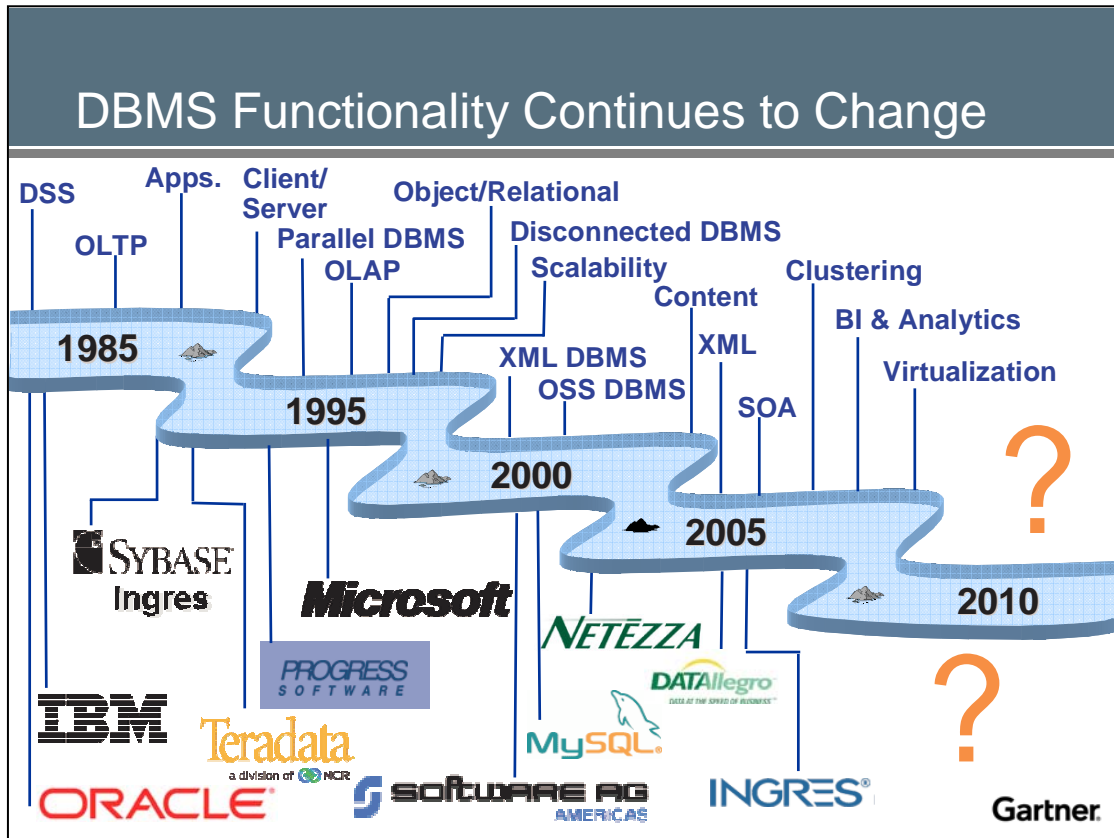
Sacramento, California

19-Oct-2006

Donald Feinberg

VP & Distinguished Analyst

Strategic Planning Assumption: Through 2010, the predominant model for new, emerging information management technologies will be via extensions to and evolutions of the original RDBMS model, including the OSS DBMSs (0.8 probability).



The original model of the RDBMS in its pure form from the 1980s no longer exists. These DBMS engines have been transformed over the past 20+ years to handle all of the new ideas and innovations that have emerged in the DBMS world. As new engines are created (e.g., MySQL), they also have a very modified RDBMS design allowing for the extension of new models. The most recent example of this is the inclusion of binary, tree-structured XML capabilities into all of the major DBMS engines.

BM and Oracle still lead in the DBMS market, as reported in “Worldwide DBMS market”, June 2006. In the windows world, Microsoft is the leader. Overall, Oracle has the predominant market share, at about 45%. These three DBMSs will continue their lead, and continue to transform their DBMS to future DBMS generations. The new players, such as MySQL, are gaining market share with OSS DBMSs, and Ingres (one of the original RDBMS engines) is making a comeback as an OSS DBMS.

What does the future hold? It's changing fast. Linux is gaining strength as a DBMS platform. The fourth workload for DW is the inclusion of SOA-type services for BI and Analytics as services in the DBMS. Self-management of the DBMS engine is another area where all the vendors are spending tremendous R&D resources, trending towards zero-planned down-time DBMS environments. Today, we find all the vendors also implementing and supporting clustering and virtualization, as these environments promise lower TCO and greater flexibility for implementing and managing the DBMS platform environment.

CIOs are investing in technologies that support enterprise competitiveness

2006 CIO Technology Priorities

To what extent will your investment in each of the following technologies change in 2006?	Ranking		Spending Increase
	2006	2005	
Business Intelligence (BI)	1	2	+ 4.8%
Security enhancement tools	2	1	+ 4.5%
Mobile workforce applications	3	3	+ 3.9%
Collaboration technologies	4	*	+ 3.6%
Customer sales and service technologies	5	8	+ 3.4%
Service oriented applications and architecture (SOA, SOBA)	6	11	+ 3.2%
Workflow management	7	4	+ 3.2%
Networking, Voice and Data Communications	8	7	+ 3.0%
Virtualization (Storage, Computing, Data Center)	9	10	+ 2.9%
Legacy application modernization and upgrade	10	5	+ 2.5%

* New question for 2006

** New question for 2005

Gartner

KEY MESSAGE: CIOs are looking to invest in technology tools that can help them deliver information and projects that will help make their business competitively different rather than commodity technology solutions.

CIOs see technology tools, rather than technology solutions as important to supporting growth and competitive advantage.

CIO views on technology priorities and spending remain focused on increasing the level of information and security across the enterprise. The top three priorities reflect the need to devote resources to

Business Intelligence – improving the collection, consolidation, and consistency of information gathered across the enterprise.

Security Enhancement Tools – keeping the enterprise, its information and brand secure.

Mobile Workforce Applications – extending central office capabilities into the field and the point of customer need.

Increased emphasis on collaboration and customer-facing technologies reflects the increased importance of business expectations for supporting growth and competitive advantage.

A view on these priorities is that CIOs are looking to invest in ways that can help them deliver information and projects that will help make their enterprise competitively different, rather than commodity technology solutions.

Key Issues

1. What future strategies will prevail in the DBMS market?
2. How will BI growth affect the DBMS Market?
3. How will open source affect the DBMS market?
4. Which vendors and products will prevail in the future?

Gartner

The DBMS market continues to be interesting in several areas: the Linux vs. Unix battle; functionality of the RDBMS in areas of disconnected DBMS, unstructured data and scalability; data warehousing and business intelligence (BI) tools support; manageability and support resource reduction; open-source software (OSS) DBMS engines (for example, Ingres and MySQL); and the re-platforming of aging, pre-relational applications.

Another interesting statistic is that the DBMS market grew (in new license revenue) 7.3 percent, from \$15.2 billion in 2004 to \$16.3 billion in 2005. The RDBMS market grew 8.3 percent from \$12.8 billion in 2004 to \$13.8 billion in 2005. This is greater growth than the overall IT budget increase worldwide for the same time frame.

New players continue to enter (for example DATAlegro, Greenplum, Netezza and more) and exit through attrition or move to OSS (for example, Informix, Ingres, Red Brick, MaxDB and Supra). This trend will not change in the near future, thus giving many new options to clients as well as the headaches that accompany such vendor churn.


Bottom Line: Companies are still buying new and additional DBMS engine licenses, and this trend will continue during our visionary horizon of five years.

Key Issue: What future strategies will prevail in the DBMS market?

Strategic Imperative: Begin to consider new ways to store and distribute data throughout the applications, products and devices that use modern technology as a means to increase competitiveness.

Storage Changes

- Grid — the network becomes the database?
- Queues
- Programmatic persistence yields to a policy scenario
- Indexing beyond the database
- Microscale persistence at the asset (e.g., RFID)



What aspects of storage repositories are changing? How is the definition changed?

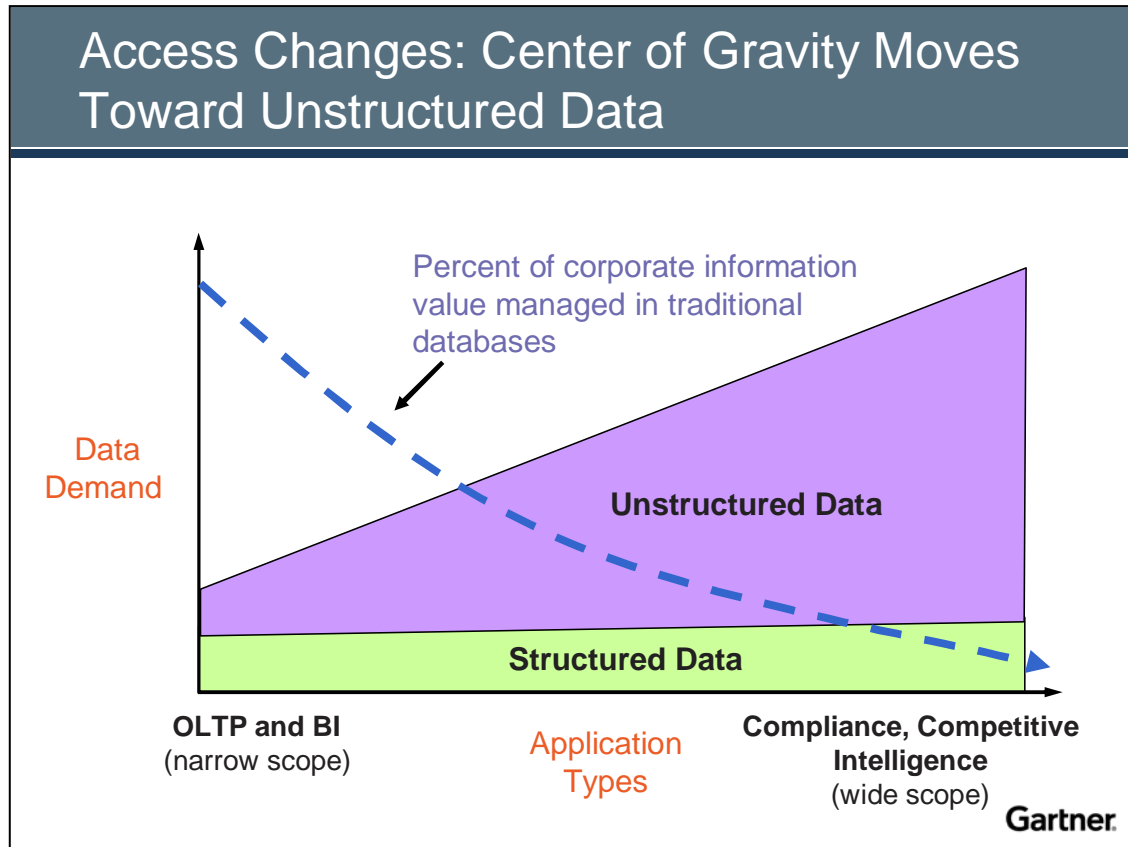
Different levels of persistence throughout the entire environment are the real nexus of changes. Current technology already creates the potential for widely distributed data storage solutions. The technology of distribution has not yet exerted even half of its influence on the traditional database system. With database vendors introducing the capability for enterprise information integration, it is a natural extension that assets start carrying identification tags with encoded descriptors and asset identification information.

There are issues with data version control powering continued needs for a centralized database. Applications that can be assembled from loosely coupled services will require a loosely coupled data strategy.

The best example of this change is the RFID chip. Is this a database? In the crate, it is a database of multiple products, and in the single-item situation, it is a database of one. In any situation it is a database and, in fact, a mobile database.

Eventually, data is everywhere, and everywhere is data.

Strategic Planning Assumption: By 2011, the traditional database will hold less than 20 percent of the data needed for most tactical decisions (0.7 probability).

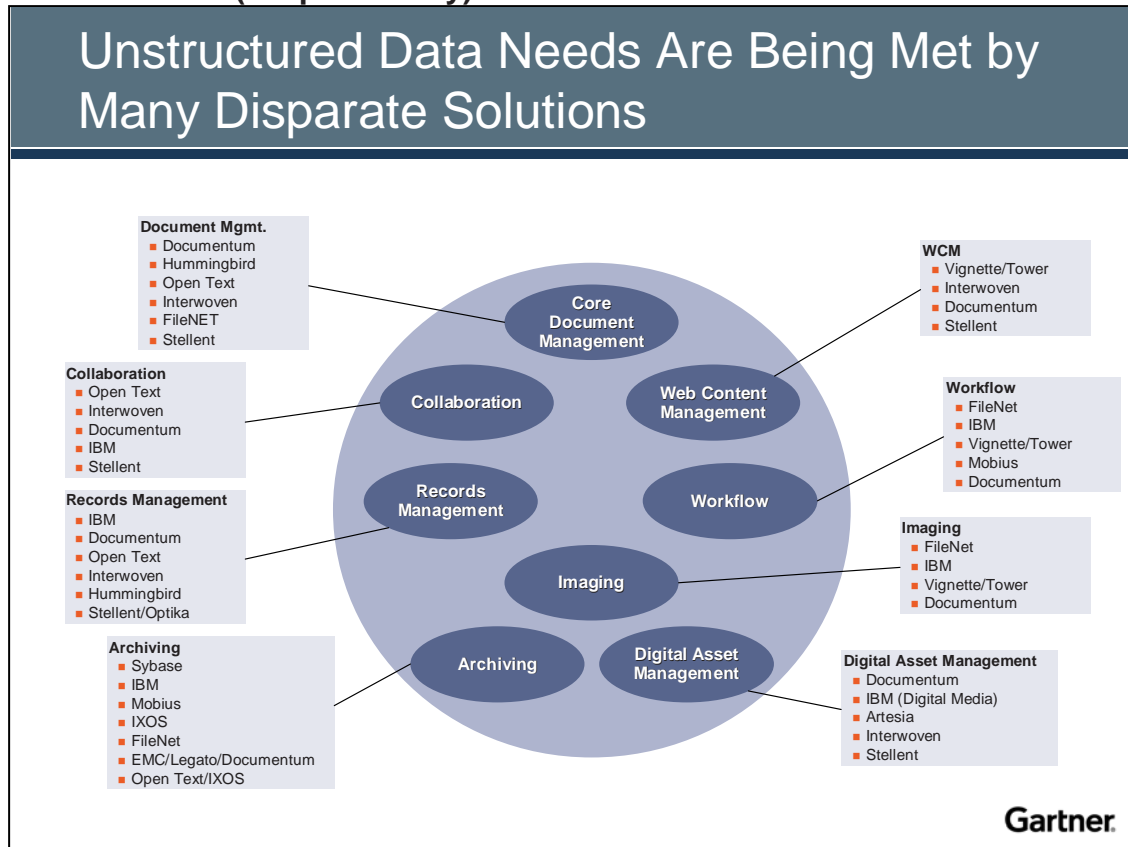


The primary issue with unstructured data is that it needs to be converted into structured data for IT to handle it. However, most unstructured data regarding process is useful for only a brief period because the state of the conditions documented are frequently transforming.

As applications evolve beyond traditional online transaction processing and BI, structured data requirements will shrink as a percentage of overall data managed by the typical enterprise. As a result, databases will bring a decreasing amount of value to the organization. New types of applications, such as compliance and competitive intelligence, will require access to and on-the-fly interpretation of unstructured data sources, such as e-mail, video and scanned document images. This data will be acquired from many sources, some of which reside permanently outside the enterprise (for example, public documents). These types of information cannot be managed well or completely in traditional databases, due in part to the limitations of DBMS products in dealing well with unstructured data types. In addition, the "completeness" of the view required for extra-organizational analytics exceeds current DBMS capability. As the perceived value of data shifts from structured to unstructured, the perceived value of traditional databases will decline, and organizations will begin to decrease their investments in resources and technology to support traditional database implementations.

Market: There currently is no true XML database market. Rather, vendors are pursuing various approaches and delivering a wide range of capabilities in this area. The depth of XML support and applicability varies greatly among these approaches.

Strategic Planning Assumption: By 2006, the major DBMS vendors will have native XML data types in the DBMS (0.8 probability).

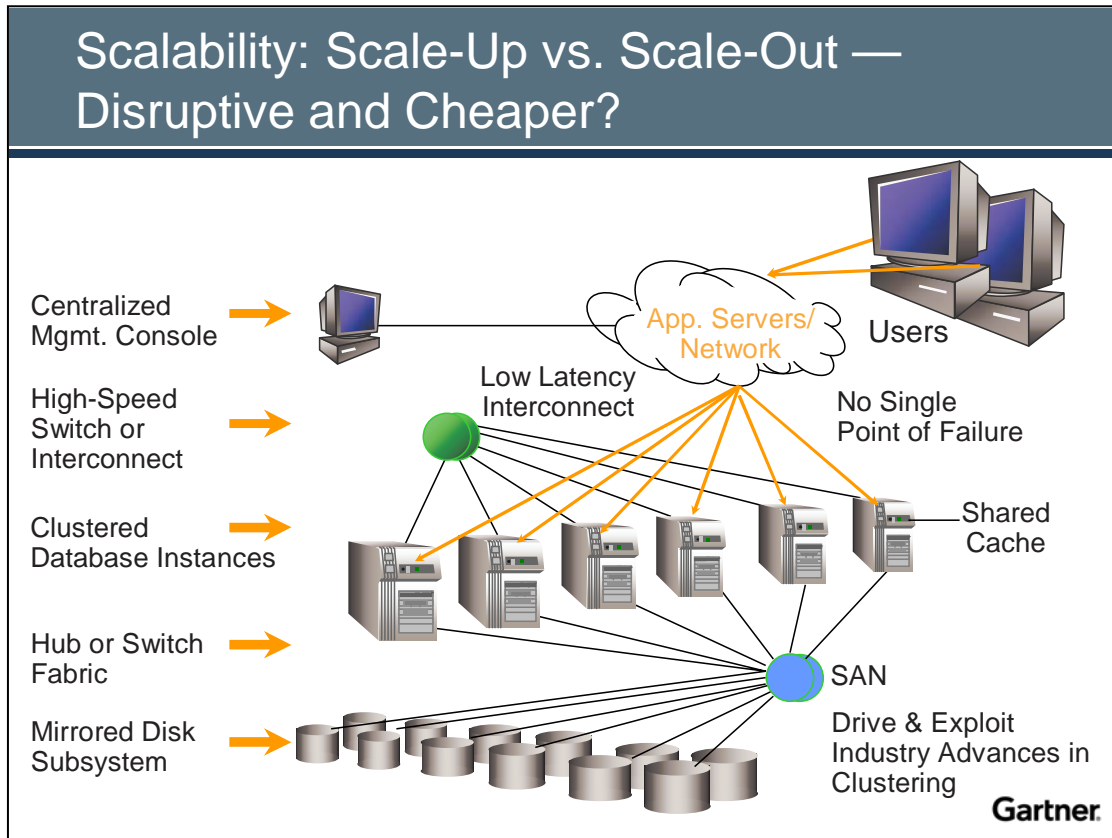


In the past few years, we've seen the new engines emerge and disappear to handle specific types of data (for example, XML). As new data types (specifically supporting the 85+ percent of unstructured data in the world) start being considered for storage in an RDBMS, we'll see the vendors of primary DBMS engines supporting this data in different ways. This includes data such as pictures, music, scanned records and faxes, as well as other data with meaning to specific types of applications (for example, medical records, security and biometric data, and geodetic data). Most of the application-specific vendors would like the RDBMS to "understand" the characteristics and content of the data. Furthermore, as we move to store all company data in a single data store, standard search engines must be able to use this data in searching.

The successful DBMS won't simply leave the data in external files, with metadata and links stored in the DBMS. The data itself will be stored in the database for integrity and functionality and will be treated as any other data type. There will be DBMS extensions (for example, IBM calls these Extenders) that enable the DBMS and the application using it to understand the data as well as its characteristics, meaning and structure.

Action Item: Begin to move your XML data to the DBMS engine for reliability, availability and security.

Strategic Planning Assumption: By 2010, 75% of highly competitive organizations will deploy high-availability data warehouses (a disruptive concept), with full fail-over and on-demand processing and storage capacity (0.8 probability).

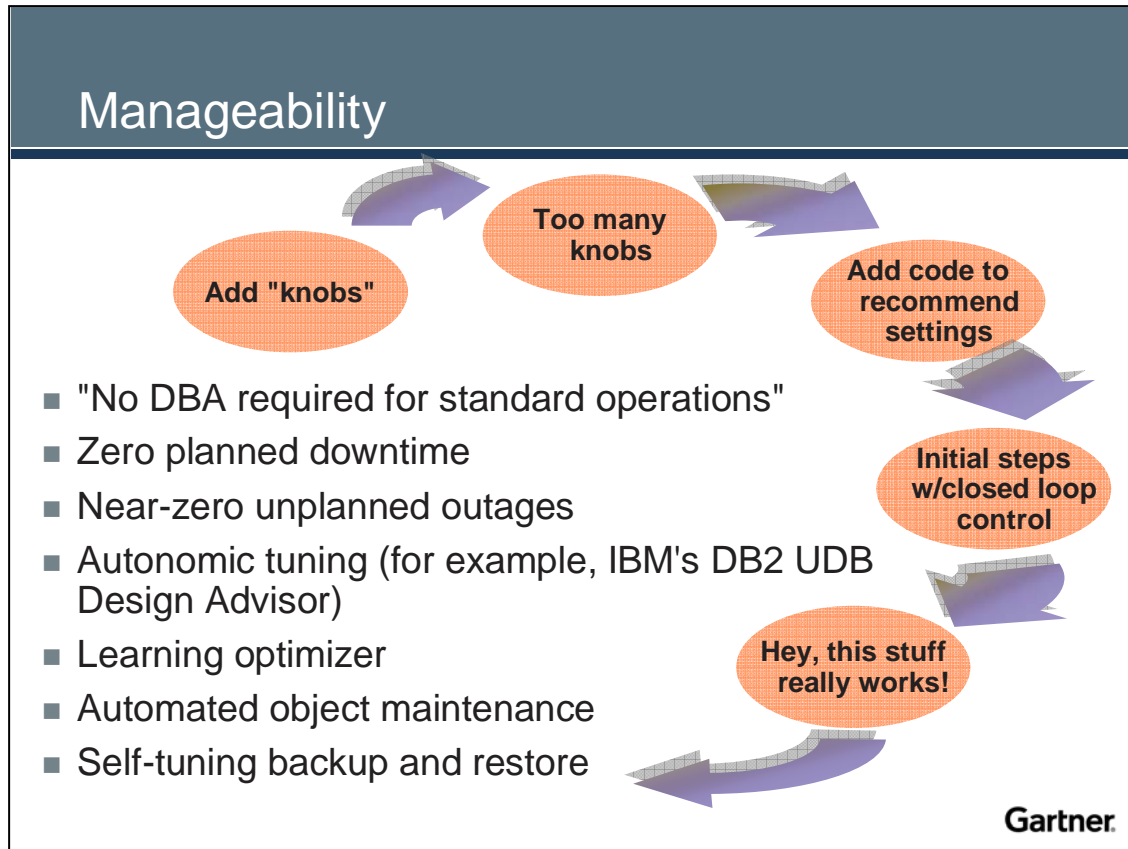


Oracle, with its Real Application Clusters (RAC) environment, originally released with Oracle Database 9i, was the first architecture to allow building systems horizontally (build-out) with commodity hardware. Today, in its fourth release, it is showing definite advantages in scaling beyond the limitations of the SMP or scale-up environments. Although originally implemented as a high-availability (HA) configuration, it is now being used in many IT environments for resource balancing and provisioning. This is especially significant when used with the new Automated Storage Management (ASM) feature released with Oracle Database 10g. The two together allow a mixed workload to be balanced at the processor level and also to be physically separated at the disk storage level. Although this would be considered a form of using data marts for performance optimization, it is still considered a best practice for many types of DW and BI applications.

Other vendors are now examining ways to allow similar scale-out architectures using processor clustering solutions and replication engines (such as supplied by the DBMS vendors or third-party software vendors such as Goldengate software). This will lead us to some sophisticated (albeit complex) configurations for scaling-out to hundreds of processors using inexpensive commodity hardware environments (i.e., lower TCO). These become necessary to support HA and mixed-workload DW environments.

Action Item: Begin to experiment with clustered solutions to gain the necessary experience to move these into production DW environments as the products gain maturity over the next several years.

Strategic Planning Assumption: By 2008, the major DBMS vendors will reach near zero percent planned downtime and reduce necessary DBA resources used to manage the DBMS by more than 50 percent (0.8 probability).



This is an area on which all vendors are concentrating. When total cost of ownership is examined, the highest cost item is people. Most companies' primary goal is to reduce the resources spent on managing DBMS engines. They're also looking for ways to leverage these resources across multiple DBMS engines.


All of the current DBMS vendors are painfully aware of these pressures. We consistently get questions about which vendor is the least expensive to install and maintain. Each of the newest releases concentrates in this area. For example, Microsoft SQL Server 2005 addresses the areas of autonomic tuning and zero planned downtime, and Sybase release 15 has also made great strides in this direction.

On the other hand, with Oracle 10g, Oracle released its Enterprise Manager with extensions into all parts of the DBMS for management. This supplies much of the functionality needed for zero planned downtime and reduction of DBA resources. Finally, the newest release of IBM's DB2 Universal Database (UDB) 8.2 also has autonomic features for zero planned downtime and minimal day-to-day intervention by DBAs. It even includes products such as the Design Analyzer, which examines a database schema and recommends changes to the DBA for better performance.

Action Item: Begin to use this functionality and allow the DBA to perform functions with higher ROI.

Key Issue: How will BI growth affect the DBMS Market?

Not All Organizations Have a Vision for BI



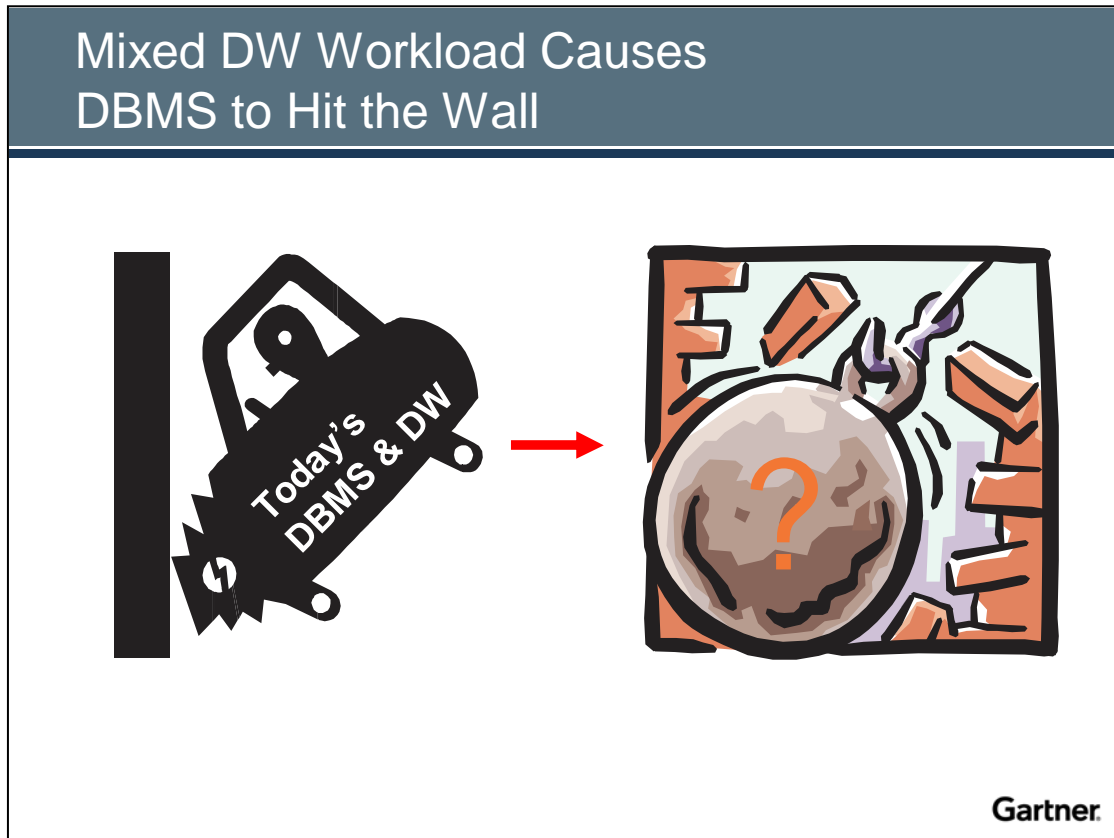
We're ignoring the need to change.

Gartner

Multiple, difficult decisions and choices must be made on the road to making BI effective and pervasive in any organization. Today, however, most of these decisions are made in isolation. In many organizations, BI is supported tactically and hasn't achieved a level of strategic importance; nor is it designed into organizations' enterprise architecture and managed as a part of their enterprise applications portfolios. The intent of this presentation is to provide clients with guidance for their BI strategic planning activities, to drive the technology and usage of BI so that it has the highest impact on business performance and transformation.

Tactical Guideline: Don't mess around with BI. Doing so will leave you deaf, dumb and blind, and affect the management and transformation of your business.


Strategic Imperative: Watching the rear-view mirror in IT will result in a significant accident —you must plan for the future and deliver now. A mixed workload today includes at least six elements; you are facing more in the immediate future.



During the past 10 to 15 years of building data warehouse environments, we have been building “giant” data marts. We define a data mart as an application-specific data warehouse. In its broadest definition, that would include today’s warehouse with reporting and/or a small amount of pre-determined queries (mistakenly called “ad-hoc”). As we push the limits of these environments, even adding sophisticated optimization layers, performance is becoming an issue. Now, as we move from batch loading to continuous loading, and from 10s or 100s of users to pervasive BI with thousands and even millions of users, we are looking to re-define the data warehouse platform.

A data warehouse platform has the capability of managing the majority of mixed workload features (continuous and batch load, operating system managed access optimization, simultaneous interoperability with multiple front-ends, supported co-location of all information types); includes management tools and techniques for manipulating the use of its entire resource base, using either physical or logical task designation; and exhibits linear scalability in all four physical aspects of data warehouse platform management.

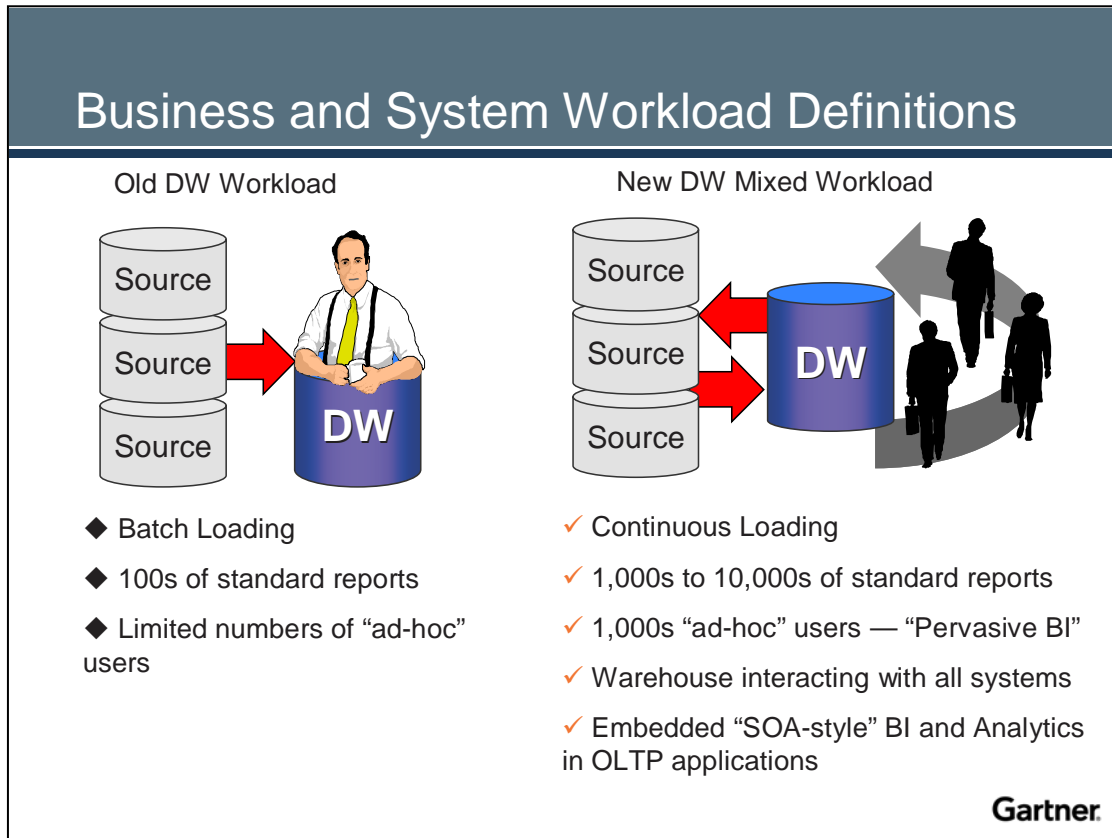
Strategic Imperative: As query complexity increases over time, the DW structure must simplify to third normal form and may require some DW appliances for performance.



The balance between query growth, query complexity and CPU resources is difficult and complex for most enterprises. Users' initial enthusiasm will drive high growth in the number of queries. However, the number of queries will decrease as interest dwindles, and the only thing remaining will be users with a real need to perform analysis. Most organizations assume that CPU resources will be adequate, but this proves to be false in most situations.

The number of queries will decrease, but users with the real requirement to perform analysis will be learning more about the data and, as a result, will ask increasingly complex questions. In addition, the decline in the number of queries is temporary in most cases. Growth returns as the deployment widens and becomes available to more business areas. The delicacy lies in keeping pace with the use of data, and in not suffering "abusive" users that drive business value but have difficulty obtaining the required resources. Organizations shouldn't permit the decline in queries to give them a false sense of security.

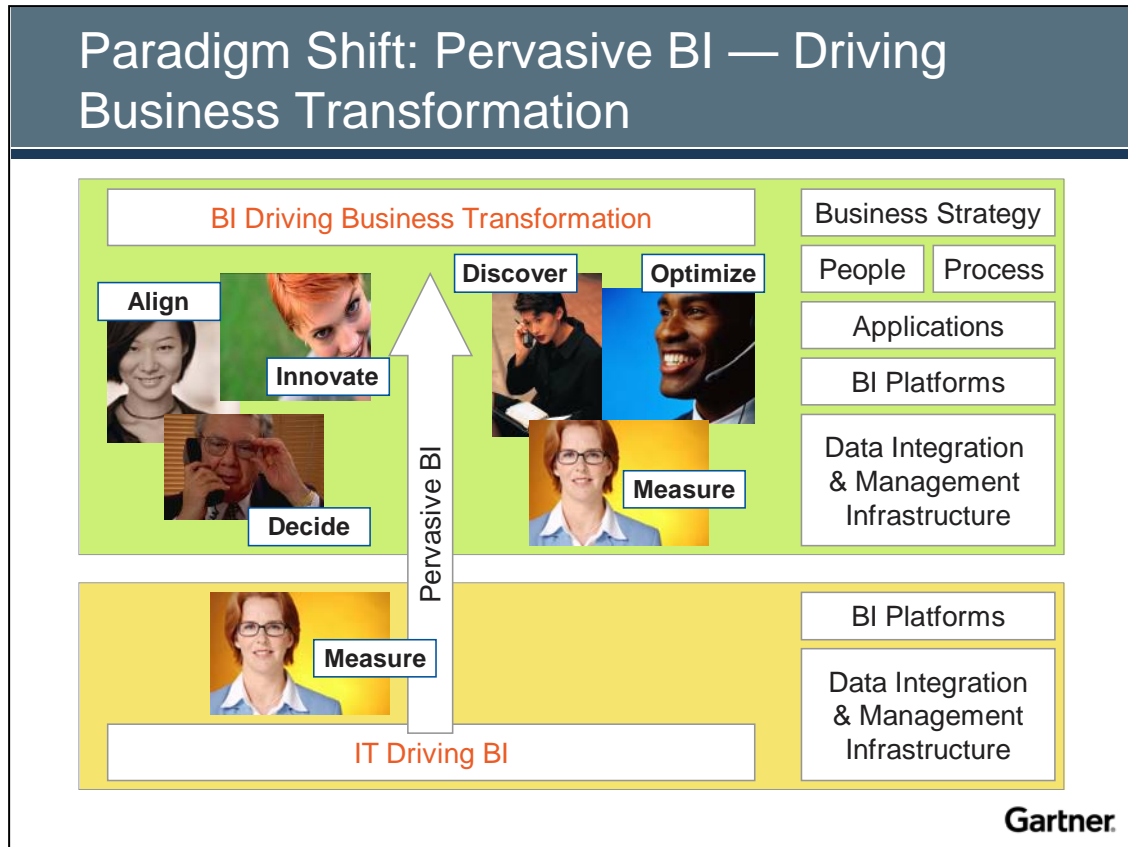
Market Definition: The DW Mixed Workload has been changing, and is becoming more complex, with a greater effect on data warehouse performance.



The traditional data warehouse and data mart worked “hand in glove” to create a usage silo for Business Intelligence. Marts and warehouses were the repository used for BI, and sometimes an Operational Data Store would also be used. With only one type of usage going against the repository, only one type of optimization was required, and only one dimension of model flexibility was needed. But the user community began to understand that there was more to using information than running canned reports and performing pre-configured drill downs (often mistakenly referred to as “ad hoc reporting”). Latency issues began to appear, not because business changed, but because business was getting more information in a better format than before. The requirement for some level of “real time” emerged from obscurity, but it was always there. This forced a different workload on the warehouse and the entire BI environment. The BI tools started to “uncouple” from the warehouse and the warehouse started to serve more applications — like BAM, CPM and others. Like the application client/server revolution before it, we are now in an information management revolution in which single perspectives on data usage are replaced with variable latency, variable optimization needs and variable end-user access approaches.

There are at least six job loads on the warehouse environment now: 1) Casual user (canned reports, filtered reports); 2) Data Miner (writes their own query, knows the model, doesn’t need you); 3) Business Analyst (advanced report writer, new drill-down report design); 4) Batch ETL (traditional, hourly, daily); 5) Continuous ETL (streaming, burst) and 6) Pervasive BI (OLTP-embedded BI).

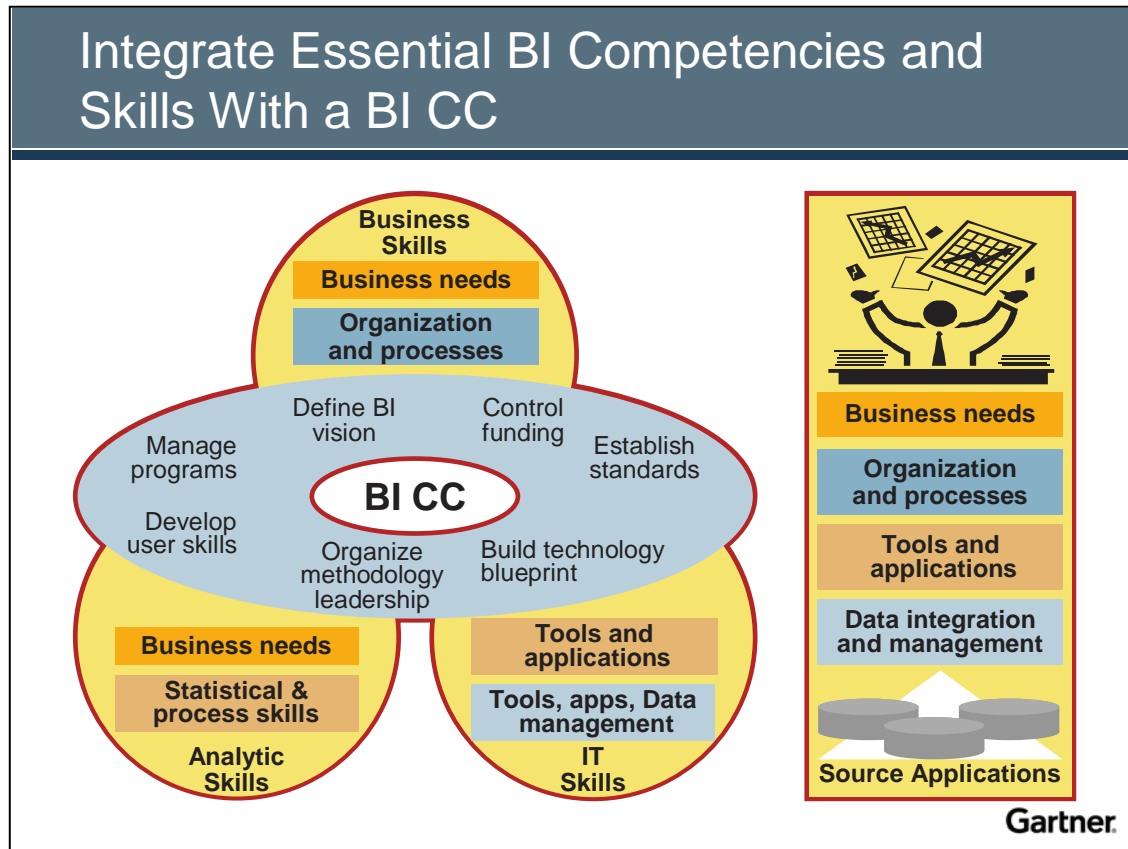
Strategic Planning Assumption: By 2012, users will interact with BI as an element of 85 percent of every business application (0.7 probability).



Business intelligence (BI) is a top initiative and investment priority for CIOs and CXOs. Many organizations view BI as "the use of information that enables organizations to best lead, decide, measure, manage and optimize performance to achieve efficiency and financial benefit." BI's value is more than information dissemination; it's highly linked to achieving business goals. We believe that BI capabilities will become more pervasive in operational and workplace applications as organizations seek to leverage BI to lead, support decisions, explore, measure, manage and optimize their businesses, and thereby drive business transformation.

Strategic Imperative: Organizations must prepare for needed changes in technology and products. More important, to use BI to drive business transformation, it will be necessary to: 1) Change the way the information architecture and application portfolio are implemented and managed; 2) change the way BI is integrated into their business processes; and 3) develop user skills and culture in the use and analysis of information as an integral part of achieving business objectives and transformation.

Strategic Imperative: A BICC that is not created with a balance of authority and power between business, technology and analytics members will not achieve their maximum potential value.



Many organizations have formed BI CCs to improve the development and focus of the resources needed to be successful with BI. These centers have been chartered to develop BI from an IT-driven initiative into a business-driven, cross-organization initiative that encompasses a wide range of users, customers and partners.

The center develops the overall strategic plan and priorities for BI, defines the requirements (including data quality and governance) and helps the organization to interpret and apply the insight to business decisions. Regardless of the organization's structure, the BI CC should report to the CIO, CFO or main business executive. In addition, it should have a mandate and a stable core, but be flexible in size.

Action Item: Assess how BI projects and skills are organized and used. Identify key resources and sponsorship toward establishing a BI CC in your business.

Key Issue: How will open source affect the DBMS market?

Linux RDBMS					
Worldwide Linux RDBMS Software Vendors for 2005, Based on Total Revenues (Millions of Dollars)					
Vendor	2004	2005	Market Share (%) 2004	Market Share (%) 2005	Growth (%) 2004-2005
Oracle	623.4	1216.5	76.1	80.6	95.1
IBM	107.7	155.6	13.1	10.3	44.5
Open-Source	31.2	45.3	3.8	3.0	45.1
Sybase	7.1	10.2	0.9	0.7	43.7
Others	50.1	81.1	6.1	5.4	61.9
Total	819.5	1508.7	100.0	100.0	84.1

Worldwide Unix RDBMS Software Vendors for 2005, Based on Total Revenues (Millions of Dollars)					
Vendor	2004	2005	Market Share (%) 2004	Market Share (%) 2005	Growth (%) 2004-2005
Oracle	4008.5	3884.8	71.3	69.9	-3.1
IBM	638.1	666.8	11.4	12.0	4.5
Open Source	21.9	31.4	0.4	0.6	43.8
Progress	117.6	123.5	2.1	2.2	5.0
NCR Teradata	390.9	417.4	7.0	7.5	6.8
Sybase	191.0	192.6	3.4	3.5	0.8
Others	251.6	245.1	4.4	4.3	-0.03
Total	5619.6	2,350.8	100.0	100.0	-1.0

Source: Gartner Dataquest
(June 2006)

Gartner

Source: Gartner Dataquest (June 2006)

Two separate issues are evolving from the open-source world: open-source DBMS engines (discussed later) and Linux. As shown above in the 2004 to 2005 Gartner Dataquest numbers, Linux is picking up speed while Unix is beginning to decline in total 2005 revenue. This trend is accelerating and will continue. By 2008, Linux DBMS revenue will be approaching 30 percent to 40 percent of the total DBMS license revenue. As Linux becomes "enterprise ready," with a proven track record, robust management tools and robust best practices, even *vendors* will completely abandon the Unix operating system. The IT organization should begin developing Linux expertise *now* and deploying Linux as an application platform wherever possible. All of the major RDBMS vendors are embracing Linux and even switching their primary development platforms to Linux (for example, Oracle beginning with 9i). Today, there are still some issues with the Linux environment and its ability to support a true production environment (*See G00140057- Data Management HC 2006*). These issues are being addressed by all of the major vendors and will diminish in the next two years.

Action Item: If not using Linux today, begin to develop expertise in Linux now to prepare for the future when it becomes enterprise-ready.

Strategic Planning Assumption: Through 2008, open-source DBMS engines will be used for non-mission-critical applications, because the other DBMS engines will deliver greater functionality, better support, better scalability, and a greater variety of management and development tools (0.8 probability).

Guidelines for Open-Source DBMS Usage (1)

- Set up a corporate governance model over the acquisition and use of OSS DBMS. Do not allow these DBMSs to be proliferated throughout the organization without control.
- Today, use OSS DBMSs for simple applications that do not require high availability and other mission-critical attributes. Watch carefully for indications of stability, scalability and reliability as they become certified by third-party vendors (e.g., applications, management tools and BI platforms).

Gartner

Today, the strongest entrant in this market is MySQL. However, considering its functional weaknesses, its rate of adoption for simple, normally "lookup"-type applications is astounding. The exception is in the IT shops, which have considerable resources at high levels to create and support these applications. IT must be careful to distinguish case studies from a top-heavy IT organization against the reality of a normal IT shop.

The four major vendors are looking for ways to put pressure on MySQL and the other OSS DBMSs by using new pricing models (for example, Oracle Standard Edition One and Microsoft Workgroup Server) and even free versions of the fully functional RDBMS engines (for example, Microsoft SQL Server Express 2005 and Sybase ASE Express for Linux). The positive effect of OSS will be the lower entry prices of the four major vendors.

Limited use of these DBMS engines — especially MySQL — is acceptable for appropriate, low-end, noncritical applications, but service from the *vendor* (not the OSS community) is recommended. If there is no clear vendor with ownership of the DBMS, do not use the engine.

Finally, take care of the marketing hype around these engines claiming market share, penetration or production usage based on number of downloads. Remember, college students download these also.

Action Item: Be sure to manage the spread of OSS DBMS within the organization by those desiring to "experiment" with these engines.

© 2006 Gartner, Inc. and/or its affiliates. All rights reserved. Reproduction and distribution of this publication in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Gartner shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.

Donald Feinberg
Sacramento, CA - 2006

Guidelines for Open-Source DBMS Usage (2)

- Use OSS DBMS engines supplied by a vendor for all production applications. If the engine is solely supported through the OSS community, it may be interesting but not for production applications in a normal IT organization.
- When using OSS DBMS engines for production applications of any type, always purchase support from the vendor supplying the DBMS. Never rely solely on the OSS community for support.
- Use the same decision criteria for choosing a DBMS platform (e.g., functionality, reliability, supportability, third-party support and stability) for an application, regardless of whether it is OSS or proprietary.

Gartner

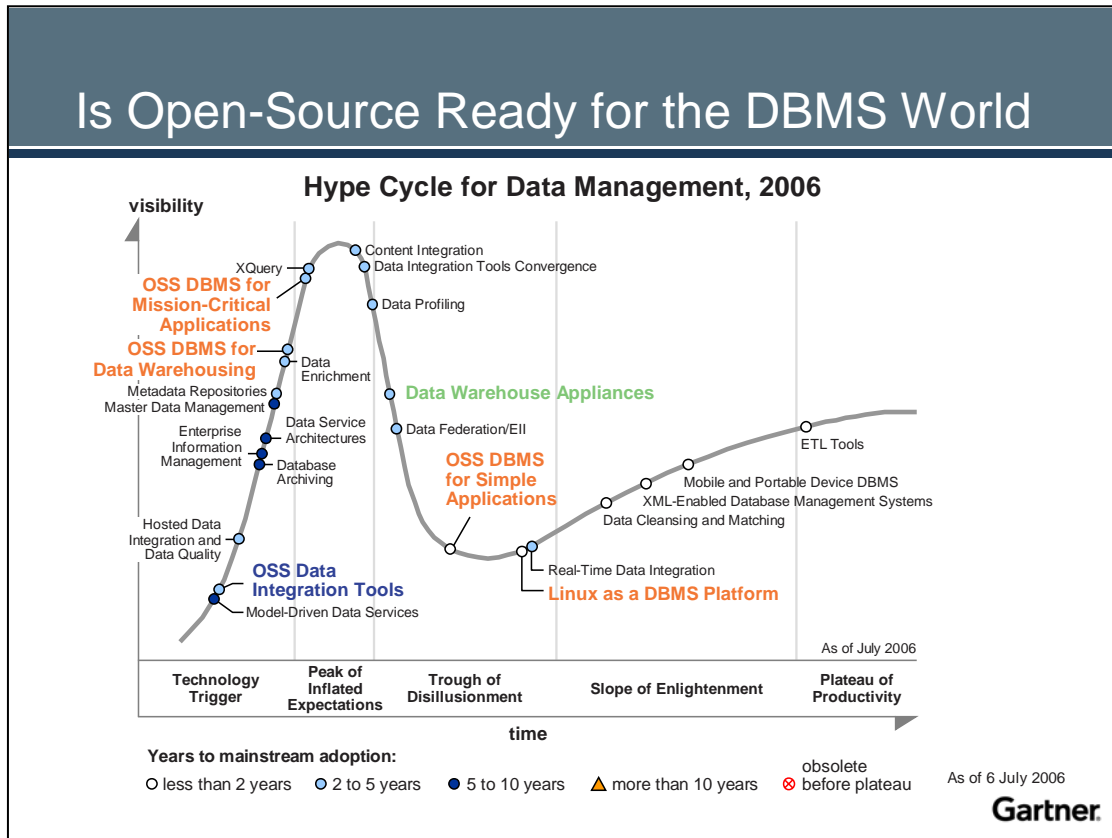
Today, the open-source DBMS market is small. Many of the entrants into this market are arguably at an "end of life" strategy. They have been placed here as means to lengthen the life span. This, along with the exposed myth of being "free," is dampening the use of these DBMS engines for anything other than noncritical applications.

With the exception of Ingres, none has a history of enterprise scalable applications and major market share. As the four major vendors continue to put pressure on this market with lower prices and greater functionality, only a few of these OSS DBMS engines will survive. Look for MySQL, Cloudscape and Ingres to have the longest life span, with MySQL the only strong new product.

Today another issue has emerged; some vendors are using these OSS DBMS engines as the basis of new engines (for example, Enterprise DB). Take care again with the marketing hype of these vendors. Basing a new engine on an OSS DBMS does not imply market acceptance, enterprise scalability or reliability.

Action Item: Use care in distinguishing useful OSS DBMS engines from the others deposited in the OSS world as a resting place.

Strategic Planning Assumptions: By 2008, open-source software (OSS) DBMSs will be used by more than 70% of IT organizations for at least one application, although few of these will be mission critical (0.7 probability). By 2010, all non-Windows-only IT organizations will be running Linux for all types of application (0.8 probability).



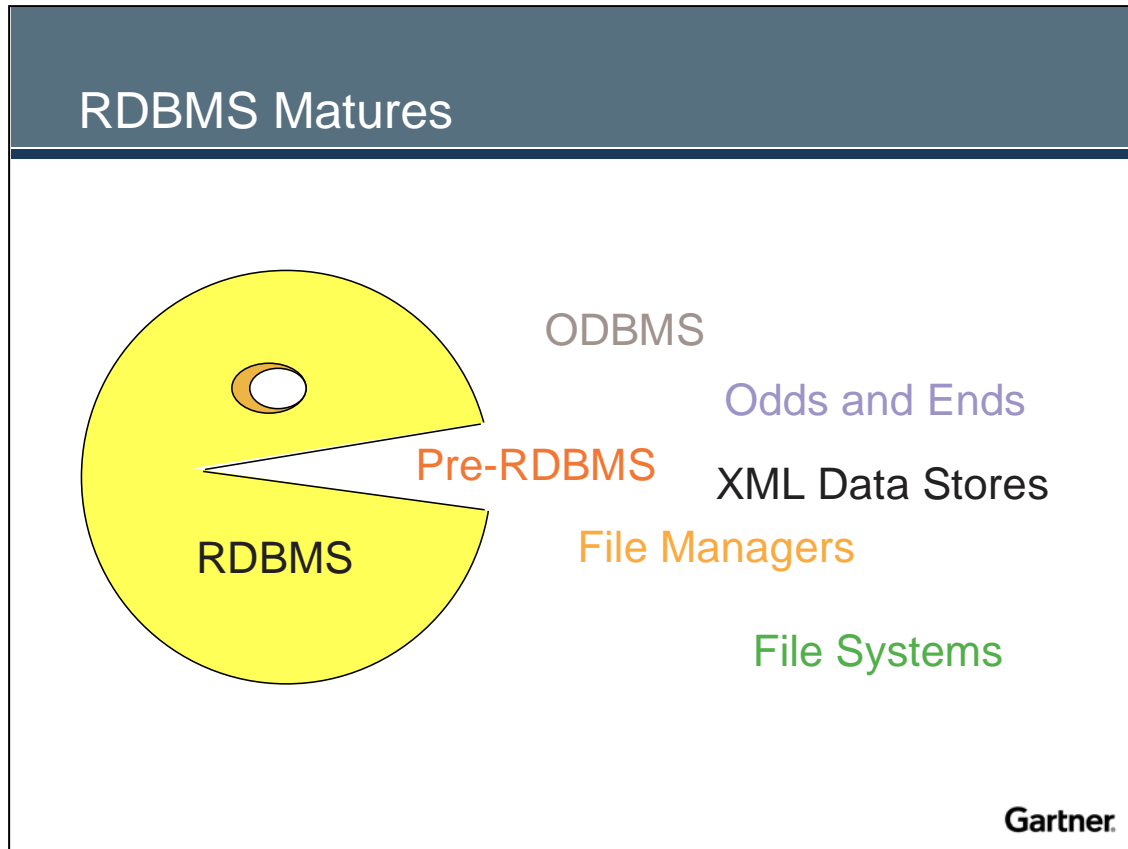
With the 2006 version of the Data Management Hype-cycle ("Hype Cycle for Data Management, 2006," G00140057), Linux as a mission-critical platform for DBMS has moved substantially and is not far into the future. During the past year, we have seen an increase in the adoption rate of Linux as a DBMS platform, primarily because of the increase in support from hardware and management tool vendors, the increasing maturity and manageability of the DBMS engines, and the increase in expertise levels of IT staff with Linux. All the DBMS vendors with Linux versions are showing increases in Linux revenue (see "Market Share: Relational Database," G00140057). The speed of adoption will continue to increase as skill levels increase, management tools mature and risk levels drop.

Open-source database servers lag behind proprietary products in terms of high-end capabilities, scalability, high availability and overall maturity. Most are not certified by major third-party software vendors, such as PeopleSoft and SAP. Furthermore, most of the tool vendors do not yet support these DBMS engines; those that do (such as business intelligence tool vendors) do so only through the standard Open Database Connectivity interface and not through native interfaces. Ingres and MySQL are furthest along the maturity line. Ingres has few third-party vendors interested. Throughout its years as a proprietary DBMS, Ingres has lost the interest of most third-party application vendors, which is a problem that will be difficult to reverse. As the new entrant, MySQL is beginning to gain interest in the tool community and among major application vendors, such as SAP.

Action Item: Begin to evaluate and use Linux for the reduction of TCO. Consider OSS DBMS engines for small, less-critical applications as a means to gain experience.

Client Issue: Which vendors and products will prevail in the future?

Strategic Imperative: Begin planning to re-platform legacy systems to the "relational" platforms; improved performance of the database systems and advances in hardware platforms will enable this move.



Several issues arise when discussing pre-relational and nonrelational DBMSs, as well as flat files. First, in the past 10 years, RDBMSs have matured and today are capable of speeds and data sizes that only pre-relational DBMSs of old could handle. Second, many of the systems are nearing the end of their 20-year life cycles, and third, there is an increasing shortage of skills for these DBMS engines. This is putting pressure on vendors of pre-relational DBMSs as more and more applications are replaced and re-platformed. Examples include CA Datacom, CA IDMS, IMS and Model 204. The Adabas DBMS on the mainframe is also of concern, but due to its availability on Linux, Unix and Windows, the pressure is less. (*See Research Note G00129706, Pre-Relational Mainframe DBMS Market Continues to Decline, 10 August 2005.*)

The second issue involves the proliferation of the "baby" DBMS in desktop systems. Examples include Access, dBase and filePro. More and more, the IT organization is being asked to take over support of thousands of these "islands of data" throughout the company. Re-platforming to small RDBMS engines is the preferred method — for example, Access is being replaced by SQL Server Express. This will continue and include Sybase ASE Express, Oracle and IBM. Finally, those who experimented with object-oriented DBMS and other "odds and ends" also are looking to re-platform to RDBMS engines.

Action Item: Develop strategies now to re-platform these DBMS applications in the next five-plus years.

Strategic Planning Assumptions: By 2007, Windows and Linux RDBMSs will scale sufficiently for small and midsize business database applications (0.8 probability). By 2009, Linux will replace Unix in large database applications and, with z/OS and Windows, will support of all large RDBMS applications (0.7 probability).

Pushing the OLTP Envelope			
High Watermarks	Windows	Linux	z/OS
Year-End 2005 Concurrent OLTP users Database size (in terabytes)	2,000 3	3,000 5	15,000 16
Year-End 2007 Concurrent OLTP users Database size (in terabytes)	5,000 10	10,000 20	30,000 30
Year-End 2010 Concurrent OLTP users Database size (in terabytes)	35,000 20	35,000 20	35,000 30

Note: High watermarks are for single SMP systems.

Gartner

High-watermark capacity depends on the technological underpinnings of a vendor's system and on the vendor's experience with installing and supporting systems of that size. A vendor's advertised "speeds and feeds" often are nothing more than self-serving vendor hype, with little or no correlation to an organization's workloads or performance requirements. Even benchmarks that include RDBMS access (a requirement for commercial suitability) are available for only a limited number of platforms — for example, Transaction Processing Performance Council Benchmark C — and shouldn't be considered a proxy for any user's specific workload. High-watermark concurrent-user levels can be attained only with great effort; "comfort zone" configurations should assume no more than 50 percent of the high-watermark figures. These figures assume a single symmetric multiprocessing (SMP) server. Clustering and other "outside the box" application-scaling techniques are implemented to meet a small minority of high-end application-scaling requirements.

Action Item: Begin to examine possibilities to replace large systems on expensive proprietary platforms with lower TCO platforms, such as Windows and Linux.

Strategic Planning Assumptions: By 2007, as Informix begins to disappear, the OLTP DBMS market will be dominated by three vendors (IBM, Microsoft and Oracle) (0.8 probability). MySQL will begin to emerge as a challenger in the DBMS market by the end of 2006 (0.7 probability).

OLTP DBMS Market					
	Strong Negative	Caution	Promising	Positive	Strong Positive
IBM's DB2 9.0					
Ingres 2006					
Informix Dynamic Serv.					
Microsoft SQL Server 2005					
MySQL 5.0					
Oracle Database 10g r.2					
Sybase Adaptive Server 15.0					

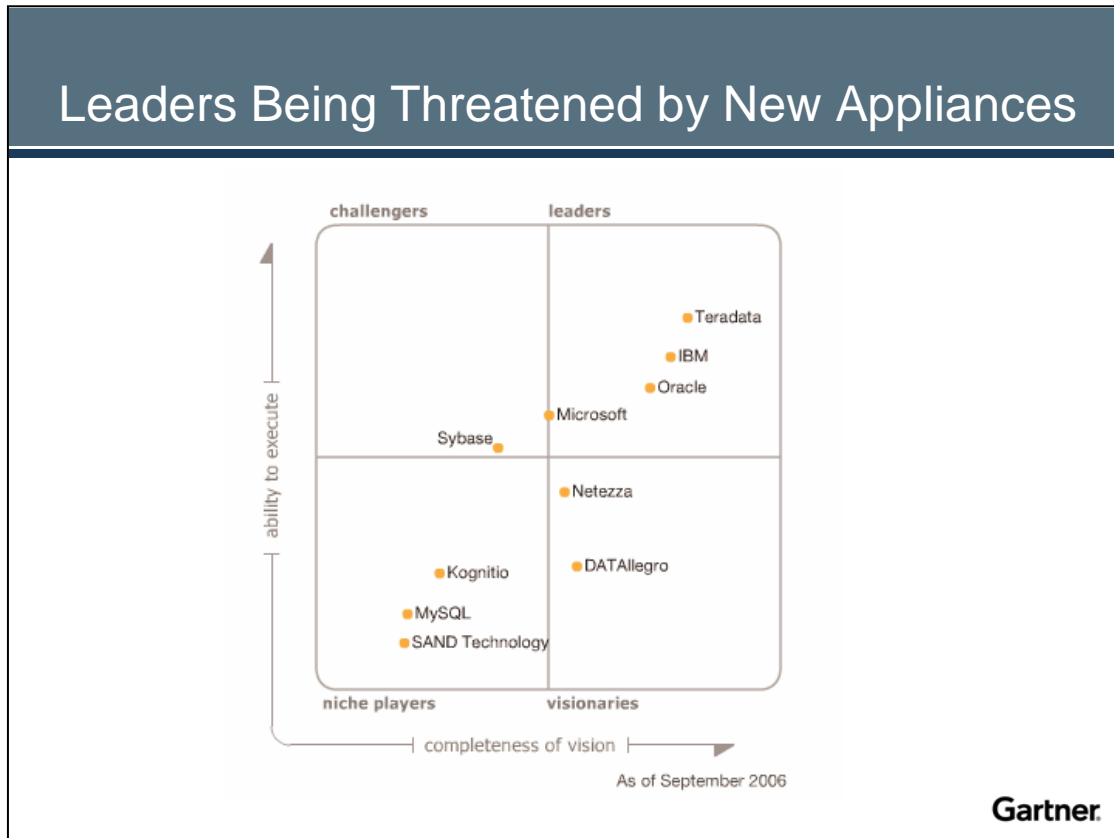
As of October 2006

Gartner

The market is still dominated by three primary vendors: IBM, Microsoft and Oracle. Sybase remains in fourth place in the OLTP market, and Teradata continues to lead the DW space. Beyond these five vendors, the only new promising entry is MySQL, which continues to gain strength in the open-source world.

IBM and Oracle continue to lead and swap positions as leader each year, with 2004 being "too close to call." Microsoft has been gaining acceptance on the Windows platform and, with the release of SQL Server 2005, will accelerate. Sybase is holding on, with new releases (r.15) continuing function feature equity, and has revenue growth largely due to its mobility offerings. If portability is an issue, Oracle leads, with Sybase in second and IBM's DB2 UDB in third. MySQL must demonstrate that it can support application vendors with functionality (not until 2008 or later) and solid customer support. This will be a major milestone for MySQL (crossing the chasm). Informix eventually will be phased out by IBM and replaced with IBM's DB2 UDB. This should be a relatively painless process for users. For most of the critical applications requiring scalability, high availability and support, IBM and Oracle remain the leaders with Microsoft's SQL Server coming on strong.

Strategic Planning Assumptions: For the DBMS through 2008, NCR Teradata will maintain its high-end DW capability lead with IBM and Oracle becoming strong players in the DW space (0.8 probability). By 2008, Microsoft will be a strong contender for the DW on the Windows platform, but still be restricted by a lack of support for other operating systems (0.7 probability).



The data warehouse (DW) database management system (DBMS) market continues to show intense competition among the leaders. During the past 18 months, some of the newer entrants have gained small to moderate market share. As IBM, Oracle and Teradata continue to battle for larger-size DWs with increased marketing and new functionality, they are faced with new competition from the appliance vendors, such as DATAlegro and Netezza, and a competitive DBMS from Microsoft. DWs are serving in an increasingly mixed workload capacity in which deep mining analysts, exploring business analysts running less-structured but equally complex queries, and fast-running tactical queries all compete for CPU, memory and disk access with differing service-level expectations. Ignore marketing claims and base your decisions on customer references and proof of concepts, to ensure that claims made by vendors will hold true in a real-life environment — more specifically, your own environment. Although this is a mature market with the full attention of large vendors seeking to make inroads with scale and innovation, smaller entrants often deliver a more focused, innovative solution.

Action Item: Continue to use tools such as the Server Evaluation Model (SEM) and Magic Quadrants to evaluate your current environment and decrease TCO, while increasing the ROI of the DW environment, especially as the workload becomes more complex.

Recommendations

- ✓ Examine applications with pre-RDBMS engines for conversion to newer DBMS engines, depending on ROI.
- ✓ Begin developing skill sets for Linux as a production platform, especially for DBMSs.
- ✓ Fence off and control open-source DBMS engines.
- ✓ Use automatic tuning functions to reduce resource consumption on day-to-day DBMS support.
- ✓ Use commodity hardware and build-out architectures to reduce costs.
- ✓ Balance vendor lock-in (stack) with new vendor innovation.
- ✓ Carefully examine your vendor's overall architecture and how it relates to your information architecture.

Gartner